NASA Conference Publication 2180

Automated Decision Making and Problem Solving

Volume I - Executive Summary

By
Ewald Heer

Proceedings of a conference held at NASA Langley Research Center Hampton, Virginia May 19-21, 1980

NASA

NASA Conference Publication 2180

Automated Decision Making and Problem Solving

Volume I - Executive Summary

By
Ewald Heer
University of Southern California
Los Angeles

Proceedings of a conference held at
NASA Langley Research Center
Hampton, Virginia
May 19-21, 1980

NASA
National Aeronautics
and Space Administration

Scientific and Technical Information Branch

1981

BLANK PAGE

PREFACE

On May 19-21, 1980, NASA Langley Research Center hosted a Conference on Automated Decision Making and Problem Solving. The purpose of the conference was to explore related topics in artificial intelligence, operations research, and control theory and, in particular, to assess existing techniques, determine trends of development, and identify potential for application in automation technology programs at NASA. The first two days consisted of formal presentations by experts in the three disciplines. The third day was a workshop in which the invited speakers and NASA personnel discussed current technology in automation and how NASA can and should interface with the academic community to advance this technology.

The conference proceedings are published in two volumes. Volume I gives a readable and coherent overview of the subject area of automated decision making and problem solving. This required interpretation, synthesizing, and summarizing, and in some cases expansion of the material presented at the conference. Volume II contains the vugraphs with various annotations extracted from videotape records and also written papers submitted by several authors. In addition, a summary of the issues discussed on the third day has been published separately in NASA Technical Memorandum 81846.

BLANK PAGE

CONTENTS

BLANK PAGE

CONTRIBUTORS

The following individuals have contributed written material to the Conference, as shown in Volume II of this report, and have given a one-hour tutorial lecture at the Conference.

Bruce Buchanan
Stanford University

Salah Elmaghraby
North Carolina State University

William Ferrell
University of Arizona

William Gevarter
National Aeronautics and Space Administration

Ewald Heer
University of Southern California

Richard Rashid
Carnegie-Mellon University

Gary Whitehouse
University of Central Florida

Charles Rieger
University of Maryland

George Saridis
Purdue University

Jason Speyer
University of Texas

Gerald Sussman
Massachusetts Institute of Technology

Robert Tenney
Massachusetts Institute of Technology

Gerald Thompson
Carnegie-Mellon University

BLANK PAGE

I. INTRODUCTION

The NASA Langley Research Center, in cooperation with the University of Southern California, organized the Conference on Automated Decision-Making and Problem Solving during the spring of 1980. The Conference was held at the NASA Langley Research Center on May 19-20, 1980. It was chaired by Alfred J. Meintel, Jr., LaRC, and coordinated by Walter W. Hankins, LaRC. The program sessions were organized and chaired by Ewald Heer, USC, under the above mentioned cooperative contract monitored by Jack E. Pennington, LaRC.

The prime purpose of the Conference was to explore topics in artificial intelligence, operations research and advanced control theory in the context of automated decision-making and problem solving as these are related to space mission oriented machine intelligence and robotics technology. It was felt that these three disciplines often require and use the same or similar techniques in solving problems, and that problem statements and avenues of research often appear different only because of differences in terminology but not in substance. The contributors to the Conference, therefore, were selected and invited based on their perceived involvement in more than one disciplinary area. The invited contributions were requested to have a slant towards overviewing the field with emphasis on assessing existing techniques, determining research trends, and identifying potential for application in NASA programs.

The contributions consisted of approximately one-hour lectures including vugraph presentation and question and answer period. The lectures were videotaped, which made it possible to annotate some of the vugraphs and extract

additional information for this report. In addition, several contributors submitted a copy of a paper for inclusion in the final report. All contributed material is included in Volume II of this report.

Volume I of this report includes a synthesis and structured presentation of the material presented at the Conference. In most cases, the material had to be condensed and summarized, while in other cases additions were necessary for completeness and to preserve continuity and coherence. As a consequence, the exact intent of the details of a particular contribution may have been compromised or misinterpreted. However, it is felt that this has had no distorting influence on the final conclusions and recommendations included in Volume I.

II. NASA PLANNING STUDIES

NASA is little more than twenty years old. The technical achievements during this time period were astonishing. Visiting planets by probes and landing men on the Moon indicated to the world that the space program can do almost anything, if enough resources are applied. In the 1970's, the NASA space program took a more utilitarian direction. The Landsat and Seasat spacecraft gathered enormous amounts of data about the Earth and its use. The Space Shuttle was developed to enable routine access to space so that space utilization and industrialization may be pushed ahead. This raises the question of economic feasibility or affordability of ventures into space. These questions must be addressed during the coming decade, and it must be demonstrated that the utilization of space is not only beneficial but also affordable.

The technologies that are expected to contribute strongly to making space affordable are the same (or similar) ones as in industry here on Earth, namely the technologies of automated machinery. Such automated devices range from simple transfer mechanisms to flexible automated manufacturing systems to intelligent robots. The frontier technologies in these areas are presently mainly associated with the intelligence portion of a system, i.e., with the computer brain where automatic decision making and problem solving is done, although this should not detract attention from the importance of the technologies associated with sensing, acting (mobility, manipulation, etc.), and man/machine interaction.

Numerous NASA studies, conferences and workshops were held during the past several years. In response to these developments, NASA/OAST redirected research

activities focusing on machine intelligence and robotics technologies. Automated decision-making and problem solving is now a major area of concern in the OAST technology program. This area concentrates on the major presently perceived technological bottleneck, namely software development and, in particular, intelligent software development.

For years now, the development of software capabilities has been outpaced by the developemnt of hardware technology of silicon chips with integrated circuits and microprocessors. The figures of merit of such computer hardware systems, such as memory storage, power efficiency, size, and cost, have been doubling approximately every year. This offers opportunities and tremendous challenges to the computer scientists and engineers to use the available hardware capabilities effectively in the development of automated decision-making and problem solving programs for the cost effective management and implementation of NASA projects.

The current NASA/OAST research and technology program for automation recognizes this deficiency in software capabilities (Fig. 1). The program objectives are concentrating first on those components leading to increased levels of machine intelligence and perception, and then on those which enable automated manipulation and robotics. The first component in Figure 1 is automated decision-making followed by relational knowledge base structuring, feature extracting and learning. This NASA program plan will undergo continuous revisions and will be updated as new planning information becomes available.

Planning activities which should be pursued immediately should include studies of the state-of-the-art across the breadth of all relevant technologies. These should include: (1) the existing technologies, (2) the available

BLANK

PAGE

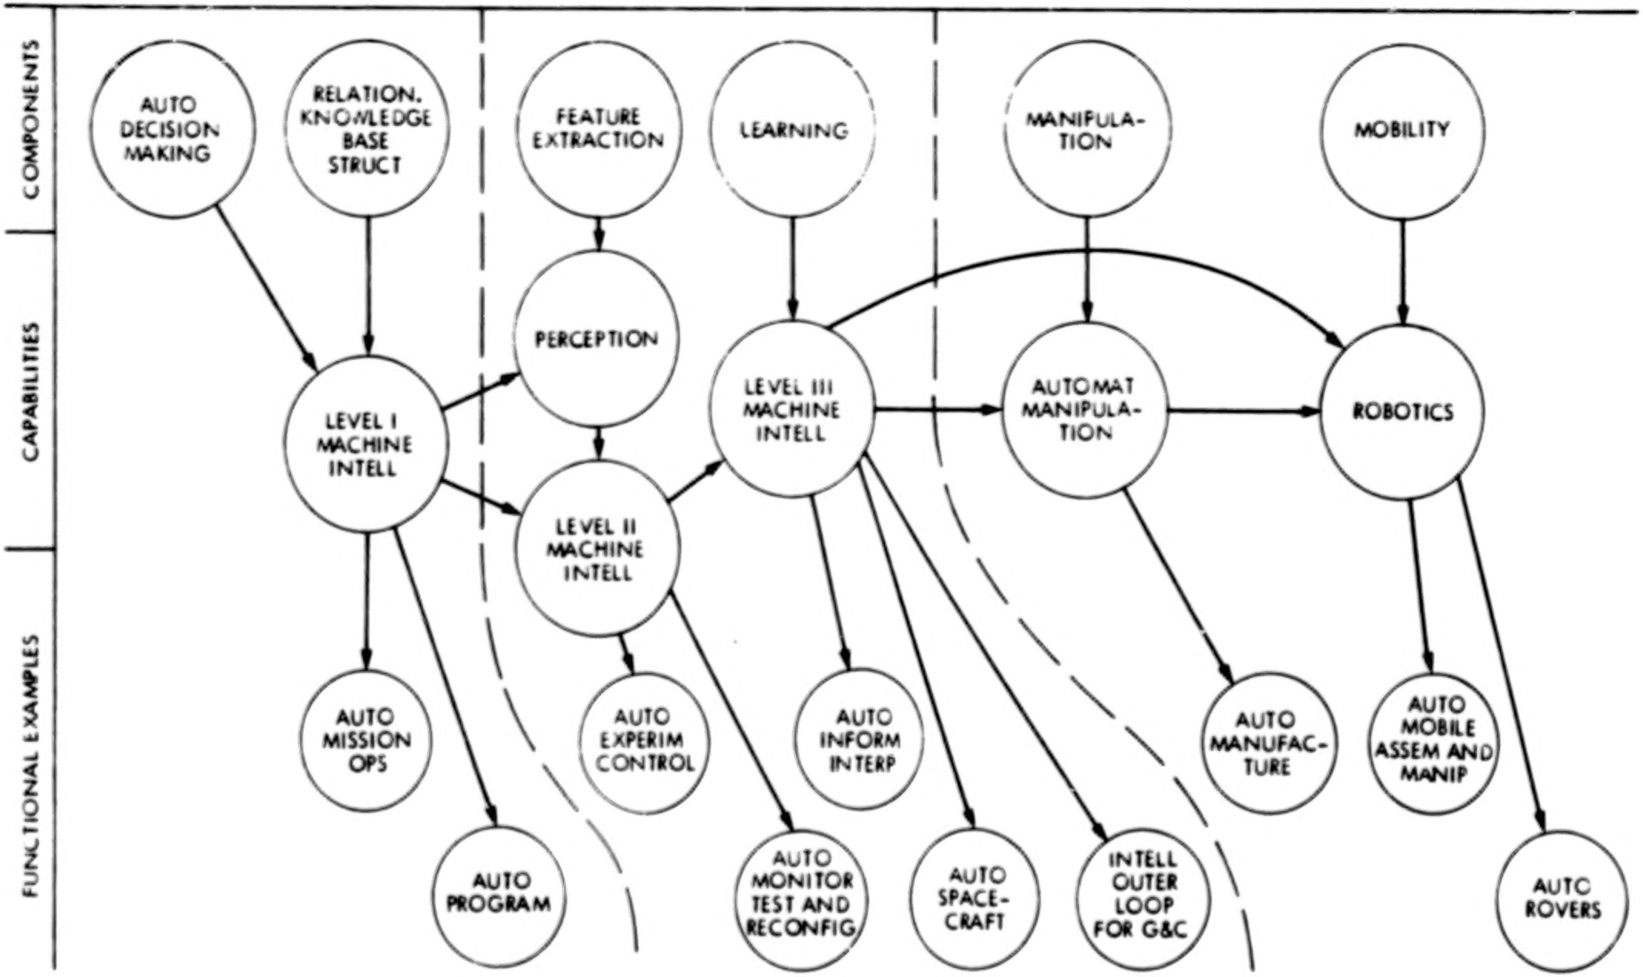

Figure 1. Automation R&T Base Program

capabilities in the country, (3) an assessment of who has these capabilities, (4) an estimate of how can NASA get access to them, and (5) future trends. From the existing modules of technology, higher levels of organizational concepts should be abstracted so that they are easier to grasp and to utilize. NASA should then selectively develop in-house capabilities in focused areas with high expected payoff. The development program, which should be coordinated with other agencies, should also pay careful attention to modularity, thus enabling portability of software modules for different applications and demonstrations of capabilities.

III. FUTURE NASA PROGRAM REQUIREMENTS

The NASA space research and development program contributes to four broad application categories:

(1) Space exploration including, among others, Earth orbital observatories and lunar and deep-space missions all of which collect data for scientific investigations.

(2) Global services consisting of a variety of Earth orbital spacecraft which collect and relay data for public service use.

(3) Utilization of space including such systems as large communication antennas, processing and manufacturing stations in Earth orbit, lunar bases, manned space stations, and satellite power systems.

(4) Transportation systems including orbital transfer vehicles, teleoperated vehicle systems, free-flying robot vehicle systems, heavy-lift launch vehicles, and the like.

In previous studies (e.g., Ref. 1) it was shown that in the space program the machine intelligence and automation technologies for data acquisition, data processing, and information extraction are the major technology drivers, and that the automatic decision-making and problem solving techniques play a major role in future technology developments. To focus on automatic decision-making and problem solving techniques is, therefore, an essential initial step towards the development of broad based automation and machine autonomous capabilities.

In the future, highly autonomous exploratory robots in space are anticipated. Such exploratory robots would communicate to Earth only when contacted or when a significant event occurs and requires immediate attention on Earth. Otherwise, they would collect the data, make appropriate decisions, archive them, and store them onboard. The robots could serve as a data bank, and their computers would be remotely operated by accessing and programming them from Earth whenever the communication link to the robot spacecraft is open. Scientists would be able to interact with the robot by remote terminal. Indeed, the concept of distributed computer systems, presently under investigation at many places, could provide to each instrument its own microcomputer, and scientists could communicate with their respective instruments. They could request special data processing onboard the spacecraft and then request that the data be communicated to them in the form desired. Alternatively, they could retrieve particular segments of raw data and perform the required manipulations in their own facilities on Earth.

Prime elements in this link between scientists and distant exploration robots would be large antenna relay stations in geosynchronous orbit. These stations would also provide data handling and archiving services, especially for inaccessible exploration robots, e.g., those leaving the solar system.

It is clear that the automated decision-making and problem solving capabilities by computers will be a major element in the operation of these systems. The robot spacecraft must be able to activate sophisticated control strategies and manage its onboard resources. It must be able to assure self-maintenance of its functions and make the appropriate decisions towards achieving the assigned tasks.

In a similar fashion, this also holds for autonomous global service robots which orbit Earth. These earth orbital robots differ from exploratory robots

primarily in the intended use of the collected data; they collect data for public service use on soil conditions, sea states, global crop conditions, weather, geology, disasters, etc. These Earth orbital robots generally acquire and process an immense amount of data; however, only a fraction of the data is of interest to the ultimate user. Most of the data could be discarded immediately since it is highly repetitive and usually is well known. Hence, the usual purpose of global service robots is to collect time-dependent data in the Earth's environment which have the character of a "new event." These data are then used to determine specific patterns or classes of characteristics, translate these into useful information, and transmit this information to the user.

Present and projected developments in machine intelligence suggest that in the future many of the presently ground-based data processing and information extraction functions can be performed onboard the robot spacecraft. Only the useful information would be sent to the ground and distributed to the users, while most of the collected data need not be retained and can be discarded immediately. This will require that the robot is able to make decisions on what data to retain and how to process them to provide the user with the desired information.

Space industrialization systems, including space utilization and space transportation systems, require a broader spectrum of robotics and automation capabilities than those for space exploration and global services. The multitude of systems and widely varying activities envisioned in space until the end of this century requires the development of space robot and automation technologies on a broad scale. It is here that robot and automation technologies, and hence automated decision-making and problem solving techniques, will have their greatest impact.

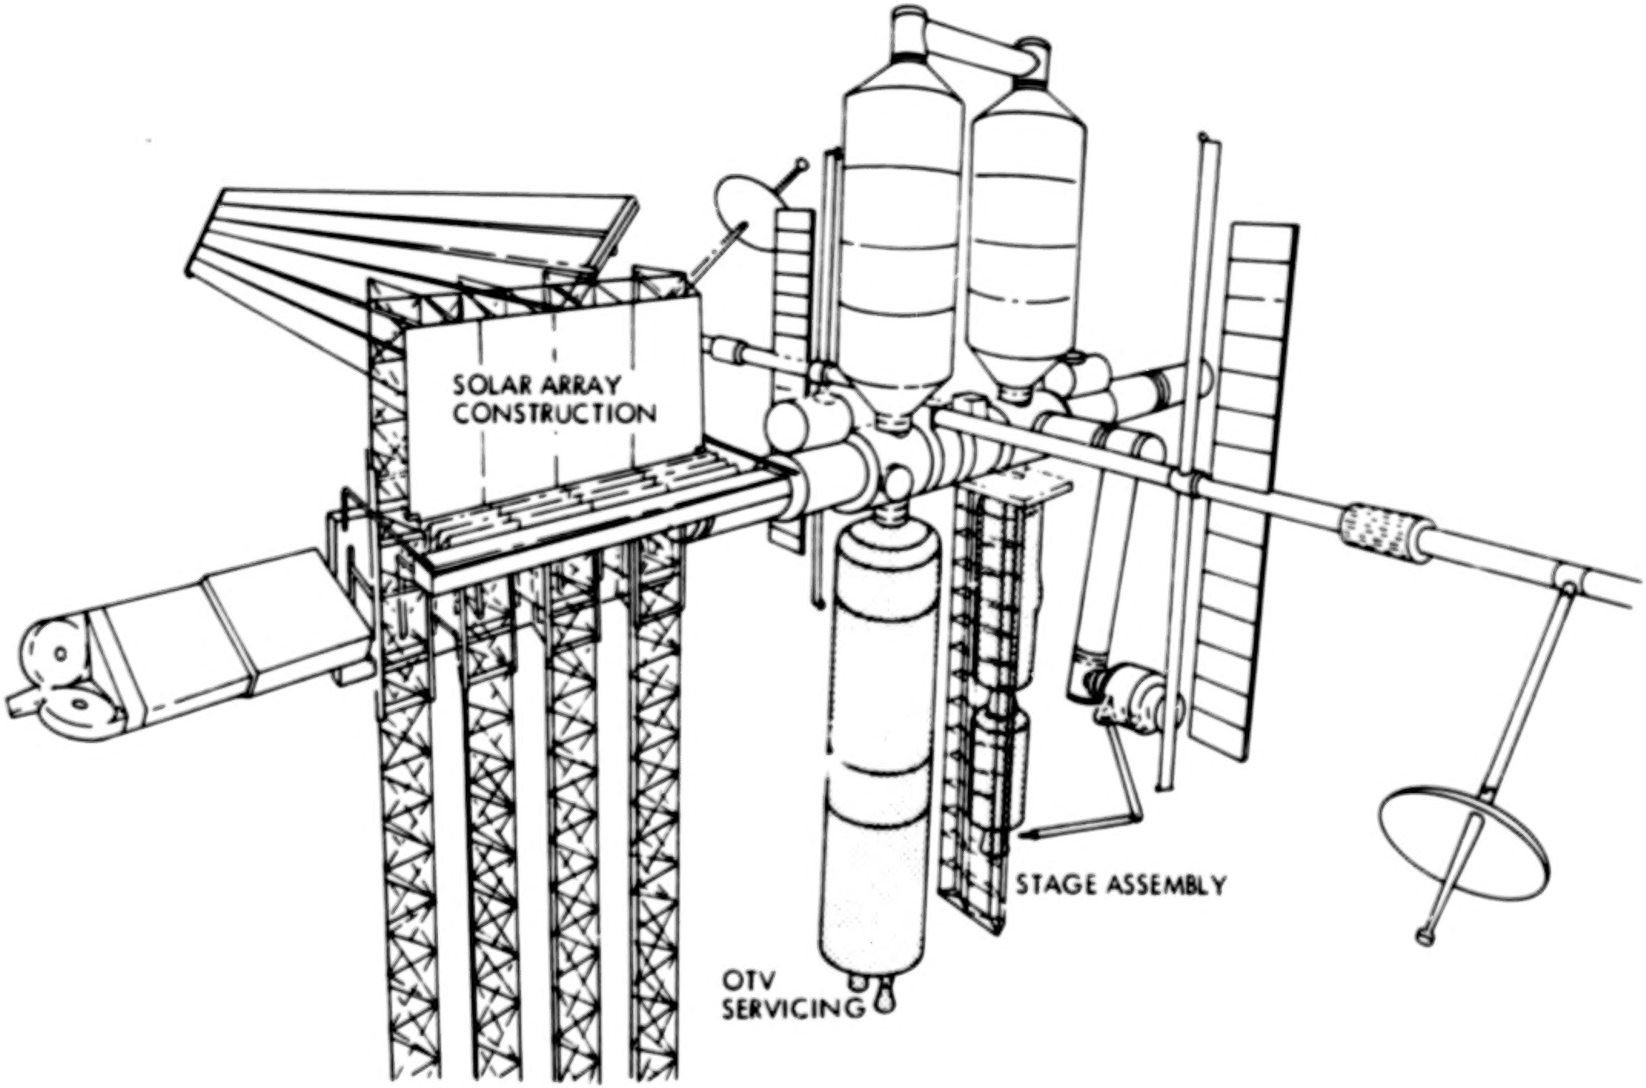

Figure 2. Space Operations Center With Construction and Flight Support Facilities

Blank Page

An example of a system, presently under study by NASA, is the Space Operations Center (SOC) (Ref. 2). The realization of SOC would put a broad based demand on the development of new techniques in advanced control theory, operations research modeling, and machine intelligence. It would require, perhaps more than any other presently planned space system, the interaction of the three disciplines. It would also require extensive investigations into the development of capabilities for the real-time allocation between humans and machines of functions which may be done either automatically or by humans.

The concept of a SOC-like system could well serve as a paradigm for the definition of space technology development requirements in artificial intelligence, operations research, and advanced control theory. As sketched out in Figure 2, the system consists of habitation modules, logistics modules, service modules, solar arrays, antennas, radiators, docking modules, and the like. In addition, there is moving machinery such as robots, teleoperators, manipulators, cranes, beam builders, and construction fixtures for building space structures. During the time of construction of the SOC and also during its subsequent operation, the structural and dynamic configuration of the system is continuously or discontinuously changing. The construction of such systems as solar arrays causes comparatively continuous dynamic changes, while docking of the Space Shuttle or the Orbital Transfer Vehicle induces strong dynamic discontinuities. The motions of the attached machines introduce both continuous and discontinuous changes, depending on the scheduling of their operations. In addition, the operating machinery may excite resonating frequencies of the system and cause damage or destruction.

It is clear that a SOC-like concept is a fertile area for identifying generic, as yet unsolved, research problems in advanced control theory, operations research and artificial intelligence as well as in areas requiring the synergistic involvement of two or all three of these disciplines. Such disciplinary interaction could develop into a broad foundation for new concepts in automated decision-making and problem solving in conjunction with human decision-making and problem solving.

In contrast to the mission application categories discussed above, automated decision-making and problem solving receive increased attention in such areas as computer aided design (CAD), computer aided manufacturing (CAM), and computer aided testing (CAT). Interactive computer aided planning is also being penetrated progressively by automation. These areas belong to the pre-mission preparatory systems engineering activities. The generic problems that must be solved, however, are in many cases the same, or similar, as for mission operations applications, and the solution techniques should often be portable from one area to another. This establishes a broad applications domain for the techniques of automated decision-making and problem solving.

IV. AUTOMATED DECISION-MAKING AND PROBLEM SOLVING

The problem solving process is inherent to the human living process. It is the process of finding ways or means towards accomplishing desired objectives, such as staying alive, providing for shelter, getting an education, driving to the airport, keeping the room's temperature at a certain level, building a bridge, sending a man to the Moon and bringing him back safely. Thus, a problem is a stated or perceived desire by one or more human beings to accomplish an objective. Humans have handled problem solving and related decision-making tasks for many thousand years, mostly subconsciously, to be sure. Some things have been learned and known about problem solving and decision-making for a long time from common sense observations. It is only during the last few decades that some scientific understanding has been gained of the processes by which humans solve problems and make decisions. It has been learned that the complexity of the problem solving process, which makes its eventual outcomes (such as building a bridge, or going to the Moon and planets) so impressive, is a complexity assembled out of relatively simple interactions among a large number of very simple basic elements. It has been shown that thinking (problem solving) processes can be synthesized with computers that parallel closely the thinking processes of human subjects in a substantial number of different problem solving tasks. The range of tasks that have been studied in this way is still narrow. However, there is little doubt that in this range, at least, it is known what some of the principal processes of human thinking are, and how these processes are organized in problem solving programs.

The above statements strongly suggest that at least some processes of problem solving and decision-making can be automated using computers. At a minimum, digital computers can be used to amplify the human mental capabilities in problem solving and decision-making tasks, as is indeed already being done. The computer is proliferating in our society in management and industry to a degree unimagined only a decade ago. The computer is assuming more and more of the functions previously done by humans and is moving the boundary more and more towards increased automation. This process calls for and requires computer programs which are based on techniques of artificial intelligence, operations research, and advanced control theory.

An automatic problem solving system has the ability to find ways and means towards accomplishing objectives given to it by humans. At the highest level, the automatic problem solving process can be subdivided into two phases which include the processes of planning and execution. At this level of abstraction, it appears appropriate here to discuss some aspects of how decision-making processes are pervading and are an integral part of the more general area of problem solving.

In general, planning can be defined as a goal oriented process of preparing a set of decisions from alternative options for action in the future. With the appropriate interpretation, this definition also holds for automatic planning, where the overall goal of a problem is stated by humans, and the computer searches for solution paths by successively decomposing the problem into subproblems, each one having its own subgoal. At each stage, there occurs a process of data acquisition through sensory input, or data base search, a

process of forecasting by modeling and extrapolation, and a process of decision-making based on synthesizing, modeling, and deducting. The decomposition of the problems into subproblems continues until the solution to each subproblem at the lowest level can be supplied by a primitive operation. The resulting plan is a data structure which represents: (1) the decomposition relations between problems (goals) and their subproblems (subgoals), (2) the ordering of operator applications that will transform the system from the initial state into the goal state, and (3) the causal relationships between actions and the goal which the actions are to achieve. There may be more than one such plan (solution path) for a given problem.

The decision on which plan (if there are more than one) will be selected for execution depends on additional considerations of implementation. It depends on such factors as resource requirements, time, reliability, complexity, etc. To be sure, any exhaustive planning effort should include all these considerations and should come up with the optimum plan. In practice, however, the tools for plan optimization (and in many cases also for plan generation) are not sufficiently developed for automation. The decisions that determine which feasible plan is to be selected are still done, to a large degree, by heuristic judgment, although some heuristic evaluations already use developed computer programs.

The second phase in problem solving is the execution of a plan. It is a goal oriented process of selecting and implementing a set of decisions within given constraints. Some of the constraint considerations are mentioned in the final paragraph of this volume. In addition, it is required that the necessary operational resources be acquired and deployed, and that actions be taken on the set of prepared decisions of the selected plan.

The implementation or action process then requires continuous control of the system variables to stay within a variety of given constraints. This control is effected by sensor and perception systems (data acquisition), by decision-making systems (synthesis, modeling, deduction), and by action systems (executing the corrections). It is clear that the control process during the execution of a plan is itself a subproblem of the execution phase, having a planning phase and an execution phase.

Extensive developments have taken place in the area of the automatic control of systems during the last few decades. It is the furthest developed and most understood area among advanced control theory, operations research, and artificial intelligence.

Advanced control theory, operations research, and artificial intelligence are the three major areas under discussion in this section. Some aspects of the man-machine interface and the human component in the automatic problem solving process will be discussed only briefly in the last subsection, which also identifies and addresses important common areas of concern.

In general terms, large and complex operational systems are usually structured in a hierarchical manner. That is to say, they are divided into units which are subdivided into smaller units, which are, in turn, subdivided, and so on. The reasons for such hierarchic structuring of systems are: (1) the components at each level of the hierarchy are themselves stable entities, (2) the hierarchic systems require much less information transmission among their parts than other types of systems of the same size and complexity, and (3) the complexity of an organization, as viewed from any particular position within it, becomes almost independent of its total size.

Complex hierarchic operational systems can be subdivided into several levels. Usually, there are three main levels with more detailed and similar functional levels within each major level. Following Figure 3, the top level sets goals, conducts the overall planning function, and establishes policy based on inputs from outside the system. The middle level (policy execution level) performs system and subsystem scheduling, supervision, and control, based on policies set at the higher level. At the lowest level (procedural level), the detailed operations are performed and controlled subject to the scheduled

HIERARCHY OF OPERATIONAL SYSTEMS

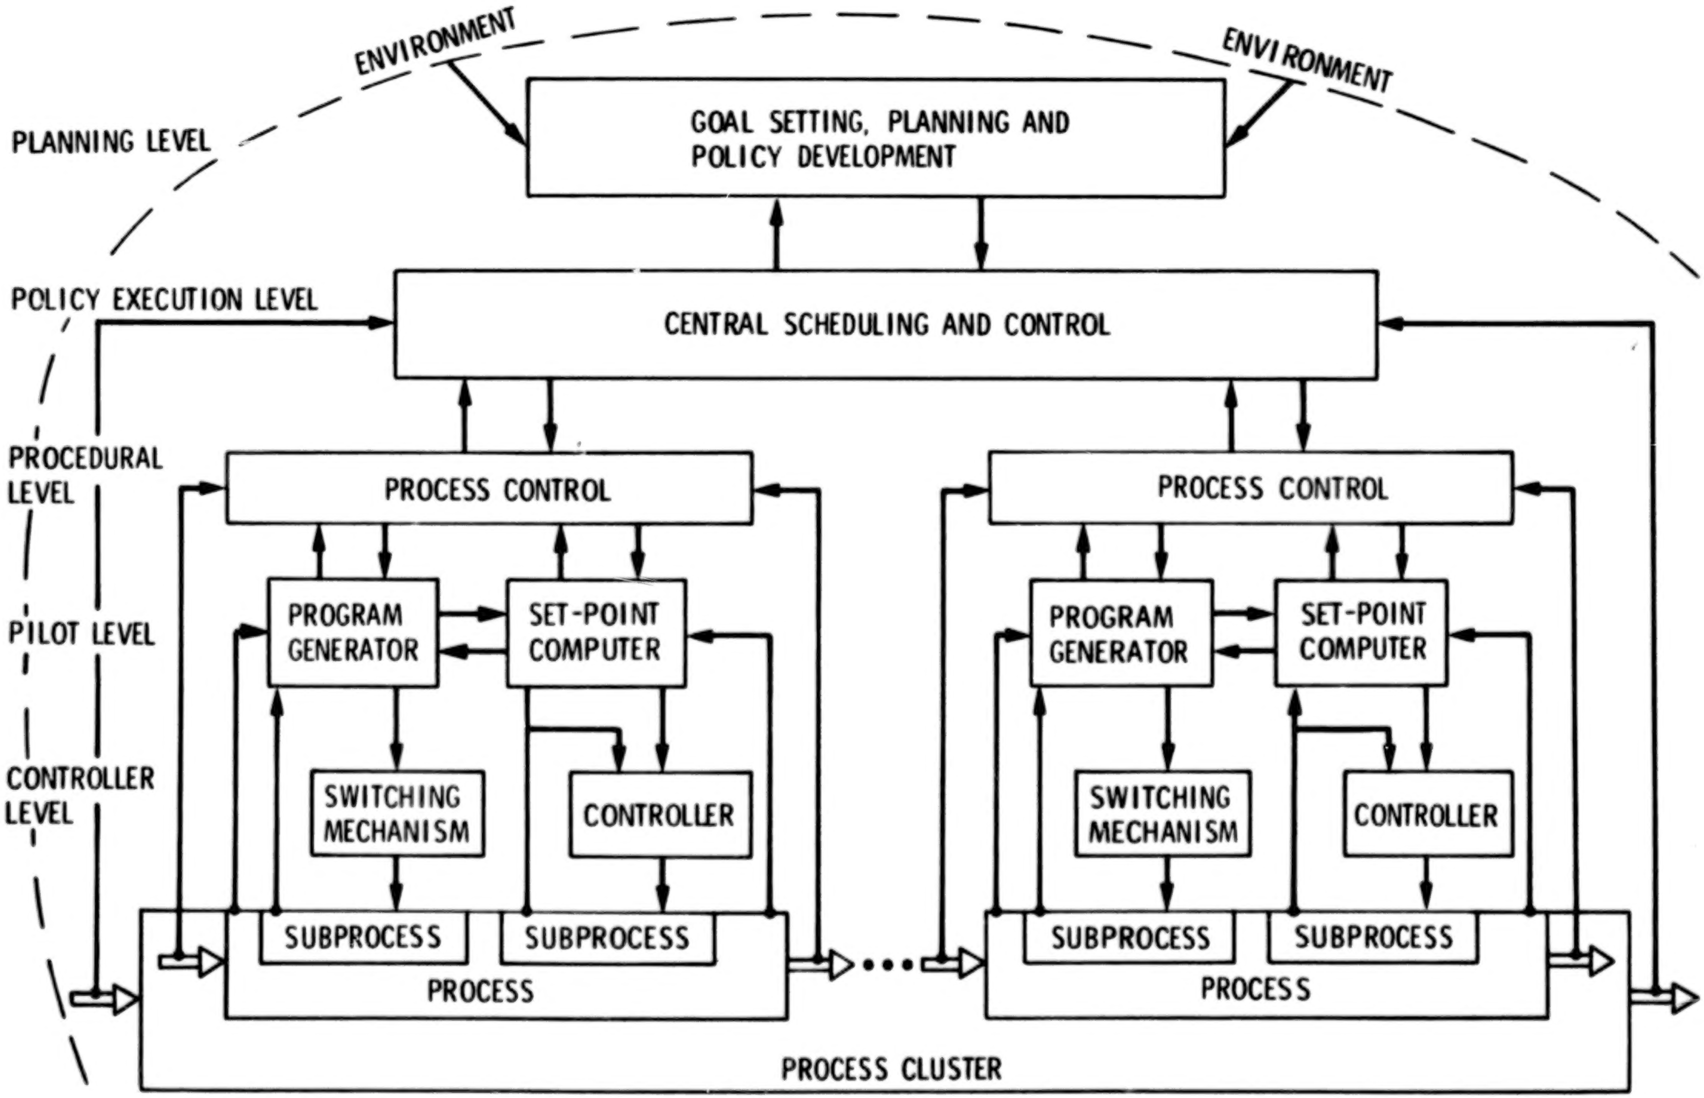

FIGURE 3

decision points and imposed control constraints. In Figure 3, the lowest level is again subdivided into three levels, where the lowest of the lowest level consists of the first level closed loop controllers and switching mechanisms, and the intermediate level of the lowest level includes set-point computers and program generators.

The problem solving and decision-making processes at the lowest level of the hierarchy of an operational system are usually to the highest degree susceptible to automation, i.e., to modeling and computer programming. As one moves into higher levels in the hierarchy, the decision-making processes become less accessible by the quantitative methods of operations research at the present state of technology. At these higher levels, it becomes necessary to use much less developed symbolic and heuristic techniques of artificial intelligence. This roughly suggests that it might be indicative to associate advanced control theory with the automation of process control (procedural level), operations research with the automation of policy execution, and artificial intelligence with the automation of policy development (planning level). However, there are strong intersections among these levels and areas, which are partly the subject of this report.

A. Advanced Control Theory

Automatic control systems provide the information flow necessary to control the movements of matter, energy and data in operational processes. The boundaries between processes and process control systems are on the one side the sensors which acquire the necessary information from the environment, and on the other side the actuators, which perform commanded actions on the environment. Between these boundary elements are the control decision making

elements which can be subdivided in order of increasing intelligence with decreasing precision into hardware control level, coordination level, and organization level.

Hardware Control Level - This is the lowest level in the control hierarchy. It usually involves the execution of certain motions and requires, besides the knowledge of the mathematical model of the process, the assignment of end conditions and performance criteria or cost functions. Optimal or approximately optimal control system theory may be used for the design of the lower level controls which belong to the decentralized subprocesses of the overall process that is to be controlled.

Coordination Level - The coordination level receives instructions from the organization level and feedback information from the process level for each subtask which is to be executed. The coordinator, usually composed of a decision-making automaton representing a context free language, may assign both the performance index and the end conditions, as well as possible penalty functions designed to avoid inaccessible areas in the space of motion. The decisions of the coordinator are obtained with the aid of a performance library and a decision-making scheme which is continuously updated to minimize the cost of the operation.

Organization Level - The organization level accepts and interprets the input commands and related feedback from the system. It defines the task to be executed and segments the task into subtasks in an appropriate sequence for execution. An appropriate subtask library and a learning scheme for continuous performance improvement provide additional intelligence to the organizer. Since

the functions of the organization level are usually performed on a medium to large size computer, appropriate translation and decision-making schemata are linguistically implementing the desirable functions.

Methodologies to determine and design the optimum control hierarchies for systems do not exist. It is, therefore, still not clear what decision-making responsibilities should reside at the various levels in a hierarchy. There are some advantages and some disadvantages that go along with centralized and distributed control processing and decision-making.

Centralized Processing - In a centralized processor system, all functions relative to problem solving and decision-making are placed in one computing system. Devices may timeshare the central processor so as to have the same effect as with a distributed processor system in which the problem solving and decision-making capability is built into the device with a small microprocessor. A centralized processor would have dynamic storage allocation built into it to provide space required by separate programs. Centralized processing has the advantage that all of the controllers or decision-makers have access to all available information. Possible disadvantages are in the memory (storage) size, processing speed, and communications bandwidth required to service the entire system.

Distributed Processing - In a distributed network of computers (micro or macro), specialized tasks can be done locally by specialized microcomputers in the network, where peripheral devices and memory in each of the processors can be shared. Distributed processors permit parallel computing to take place; however, a major advantage of distributed processing is the inherent redundancy. If one processor fails, the operations can still continue, albeit at reduced speed and reliability. The total system does not degrade its performance catastrophically; it is "fault tolerant."

Basic feedback control techniques are not applicable to problems where the system structure changes, e.g., as a result of component failures, additions to the system, wear, etc. This leads to the modern notions of control theory, which include the mathematical concepts of state, probability, and optimization, and which provide a combined framework for control decisions that are the "best" possible, but not necessarily the optimum. The formulation of control problems in these terms is about the same as the early problem solving methods of artificial intelligence. Both attacked the same question: how to manipulate the problem world to achieve desired ends. The differences in these two approaches appear through additional assumptions about the system structure. These additional assumptions enable their practical application.

Advanced control theory is best applied in contexts of systems with relatively simple structure in which either the number of states is small, or a multitude of states can be considered explicitly through symbolic manipulations. Thus, the limitations of advanced control theory becomes progressively more apparent as one has to deal with increasing complexity inherent in large and distributed systems.

The developed tools for dealing with such systems are not always suitable for practical applications. The analytical tools are often only applicable to very narrow classes of systems, and much insight is required to formulate a control problem in terms of one of these classes. The required computational resources are often not available even if the control problem can be formulated. Thus, advanced control theory can be viewed today as a library of analytical techniques, each suitable for a special class of problems. One of the most pressing issues for advanced control theory is that of dealing with questions of complexity in large scale systems.

One fairly successful approach to controlling a large scale system is to decompose it into weakly interacting, small subsystems. A second approach is to postulate the existence of several controllers, either cooperating as a team or working individually, but receiving different information about the system. Each one observes different outputs and is responsible for different inputs and system behavior.

Unfortunately, decentralization in advanced control theory has not been very successful in dealing with complexity. Even seemingly simple problems become quite complex when approached this way. Some of the reasons for this are that optimization encourages centralization, and each decision-making processor needs to consider not only the impact of its decisions on the system directly, but also their impact in the context of inputs supplied by other decision makers. This "secondguessing" is particularly bothersome when it recurses, i.e., when each of the other controllers must model the one decision-maker, which includes the latter's knowledge of the former's knowledge of the latter's, ... etc.

B. Operations Research

Operations research is a scientific approach to problem solving. The application of operations research involves the construction of mathematical, economic, and statistical descriptions or models of decision and control problems to treat situations of complexity and uncertainty. It also involves the analysis of the relationships that determine the probable future consequences of decision choices, and the creation of appropriate measures of effectiveness in order to evaluate the relative merit of alternative actions. Operations research appears to offer general techniques to deal with at least some of the

areas of complexity in automated decision-making and problem solving, namely, the effective management and coordination of many controllers, for which advanced control theory has not been successful.

Operations research offers many tools associated with names such as linear and nonlinear optimization, dynamic programming, queueing theory, combinatorial theory, network theory, scheduling theory, and the like. These tools enter the automation of problem solving and decision-making processes by providing mechanized means for performing some of the functions formerly performed by humans. Here automation begins to enter the domain where humans have retained their highest comparative advantage versus the machine, that is, the use of their brains as flexible general-purpose problem solving devices. This brings up the picture of man's functions in a man-machine system (e.g., teleoperator, Figure 4), in which some of the system's functions are performed autonomously by the machine (i.e., a robot), and the machine devices must be matched to the human characteristics.

Most operations research techniques fall into the province of programmed or well-structured problem solving. This contrasts with nonprogrammed or ill-structured problem solving for which "heuristic programming" or "artificial intelligence" techniques are used (to be discussed in Section V). There is a whole continuum between the highly well-structured decisions at one end, and the highly ill-structured decisions at the other end. In Figure 5, the two sides identify respectively some of the characteristics of well-structured and ill-structured problems for human and automation oriented decision-making and problem solving.

MISSION OPERATIONS MODEL

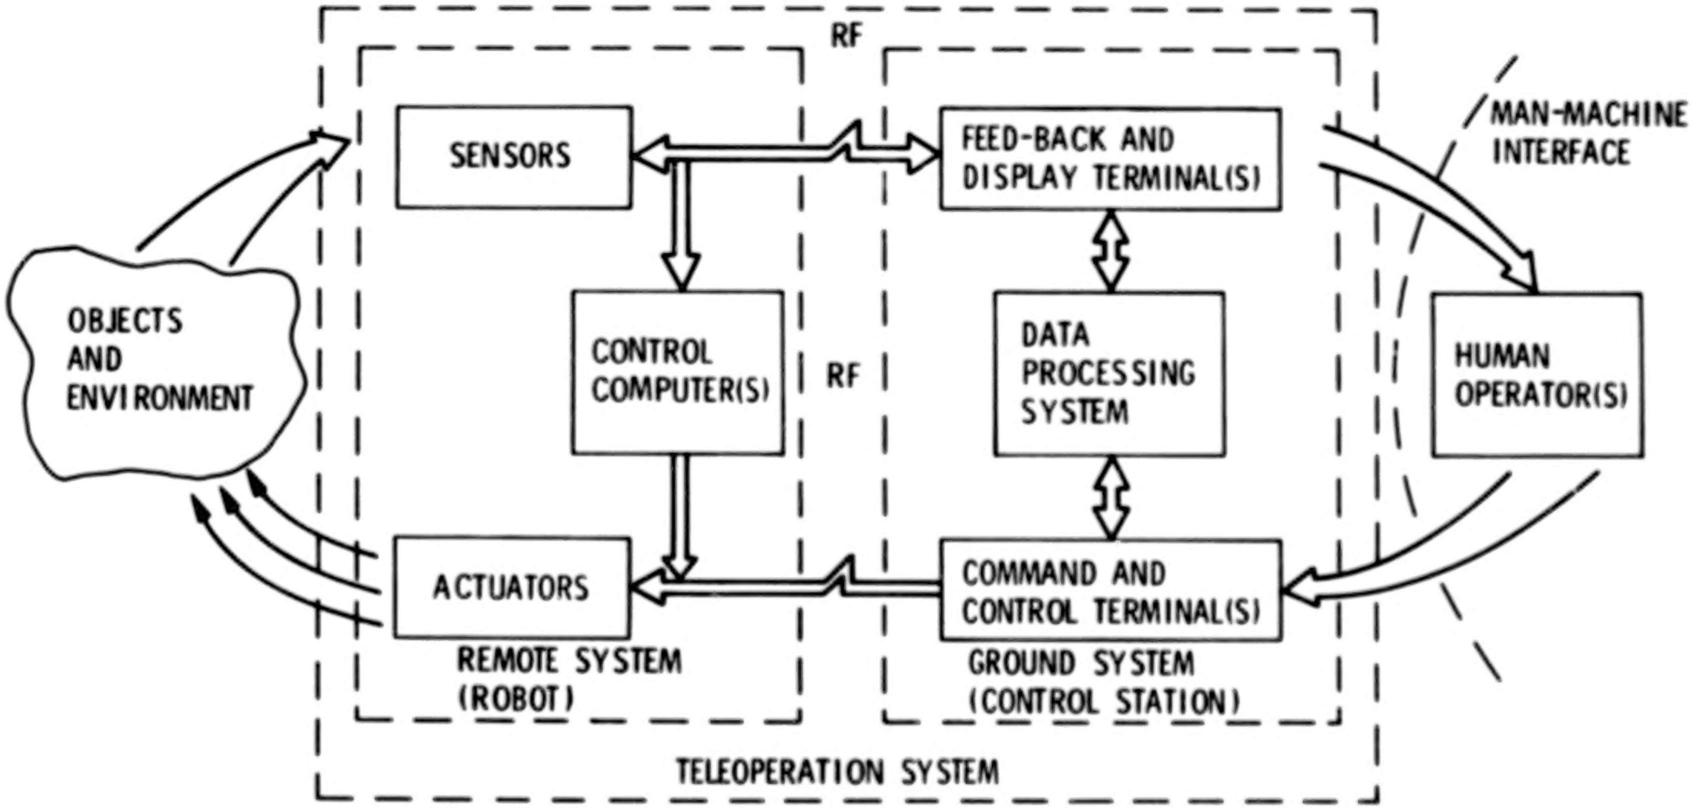

FIGURE 4

- <u>WELL-STRUCTURED PROBLEMS</u>
- <u>ILL-STRUCTURED PROBLEMS</u>

- TYPICAL DECISION TYPES

WELL-STRUCTURED PROBLEMS	ILL-STRUCTURED PROBLEMS
• ROUTINE REPETITIVE DECISIONS	• NOVEL POLICY DECISIONS
• PROGRAMMABLE DECISION PROCESSES	• NONPROGRAMMABLE DECISION PROCESSES

- HUMAN ORIENTED DECISION-MAKING METHODS

WELL-STRUCTURED PROBLEMS	ILL-STRUCTURED PROBLEMS
• HABIT	• JUDGMENT
• CLERICAL ROUTINE	• INTUITION AND CREATIVITY
• STANDARD PROCEDURES	• RULES OF THUMB
• WELL DEFINED COMMUNICATION CHANNELS	• SELECTION AND TRAINING OF MANAGERS

- AUTOMATION-ORIENTED DECISION-MAKING METHODS

WELL-STRUCTURED PROBLEMS	ILL-STRUCTURED PROBLEMS
• OPERATIONS RESEARCH	• HEURISTIC PROBLEM SOLVING
• COMPUTER DATA ANALYSIS AND PROCESSING	• HEURISTIC COMPUTER PROGRAMS

Fig. 5. Problem Structure and Solution Techniques

Whatever the specific operations research technique, the general procedure for using it in decision-making is something like this:

(1) Construct a mathematical model that satisfies the conditions of the tool to be used and which, at the same time, mirrors the important factors in the particular problem situation to be analyzed. The basic structure of the tool must fit the basic structure of the problem with occasional compromises on both sides to fit them to each other.

(2) Define a criterion function which is a measure for comparing alternative merits of various impossible courses of action.

(3) Obtain empirical estimates of the numerical parameters in the model that specify the particular, concrete situation.

(4) Carry out the mathematical calculations required to find the course of action which, for the specified parameter values, maximizes the criterion function.

Following this procedure, a program is constructed which arrives automatically at decisions to be executed by the system. In this process, some decisions that have been judgmental (ill-structured) before are being incorporated into the area of programmed decisions, while others may still remain judgmental, to be made by humans. The latter still belong to the nonprogrammed decisions. Similarly, the first three steps of the four-step procedure above fall into the category of ill-structured problems at the present state of technology. However, important advances are being made in the area of artificial intelligence to incorporate more and more of the first three steps into the well-structured domain.

An important area of operations research and systems theory is interpretive structural modeling. It has its basis in graph theory, set theory, mathematical logic, and matrix theory. Some of the most fundamental theory of structure is found in network theory. Network theory forms the basis for a wide range of different systems analysis tools and modeling techniques. The power of these tools and techniques stems in part from their flexibility and ability to represent widely different entities, such as the physical relationships present in electrical networks or machine systems, as well as many logical relationships, such as those present in computer flowcharts, scheduling, or sequencing problems.

During the last 25 years, the developments in network theory spawned the well known CPM and PERT techniques. More advanced techniques such as the GERT (Graphical Evaluation and Review Technique) were developed during the last decade. GERT combines the disciplines of flowgraph theory, moment generating functions, and PERT to obtain solutions to stochastic decision problems. This and similar decision networks differ from PERT in that each node in a decision network may have several possibilities to proceed. This also includes the possibility to return to an earlier state, and start over again (Figure 6).

The versatility of network modeling techniques has been shown in applications to transportation and resource allocation problems. These problems are among the easiest, since they are polynomially bounded; that is, in the worst case, the maximum number of steps required to solve such problems can be constrained by a bound which is a polynomial function of the amount of input data needed.

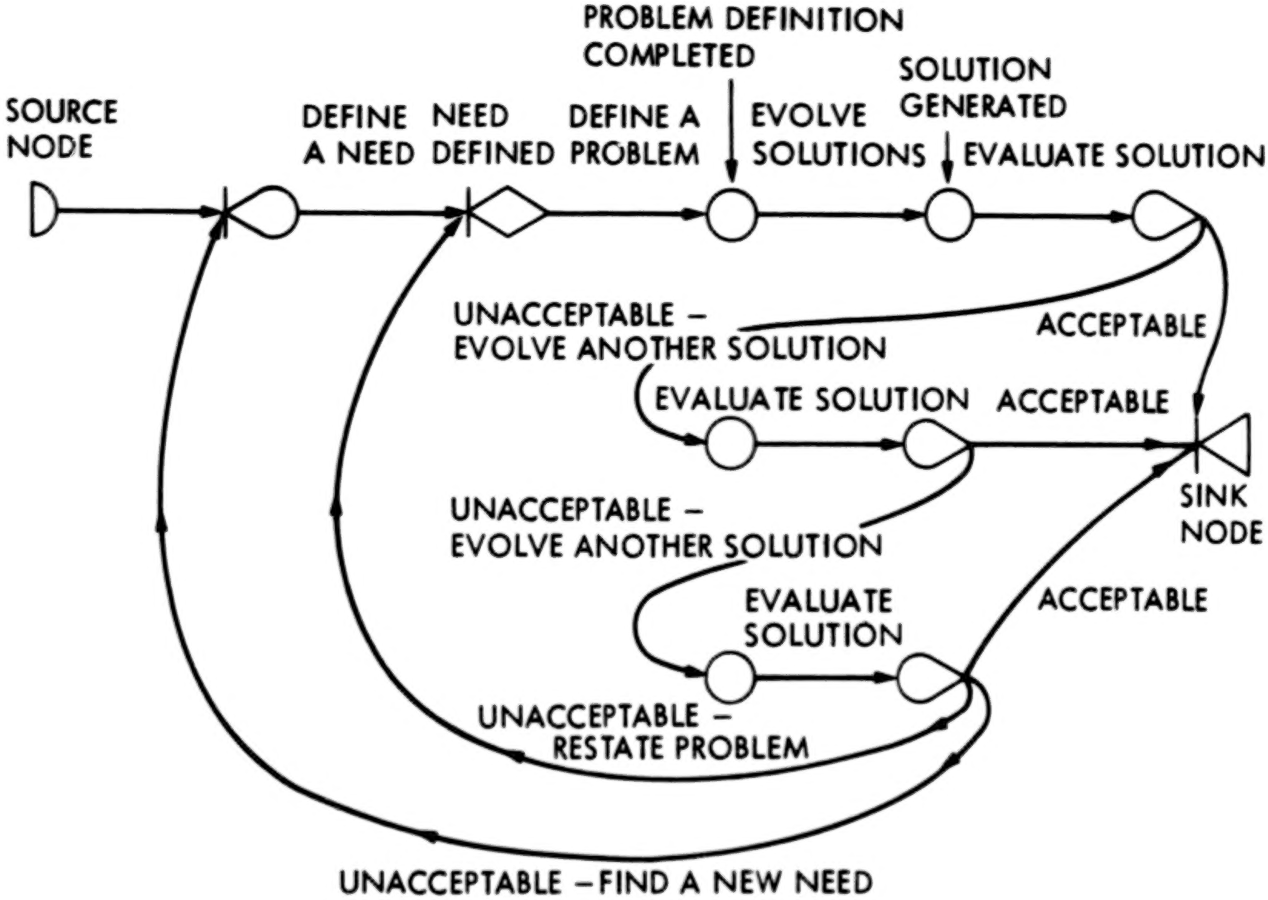

Figure 6. GERT Research Model

Since network problems usually have only a finite number of feasible solutions, it is natural to consider the use of some kind of enumeration procedure to find an optimal solution. Unfortunately, this finite number can be, and usually is, very large. For example, if there are 10 variables and each one has 10 feasible values, then there can be as many as 10^{10} feasible solutions. Despite the speed of modern computers, exhaustive enumeration would be prohibitively time consuming even for relatively small problems. Therefore enumeration procedures must be cleverly structured so that only a small fraction of the feasible solutions need to be examined.

One such approach is provided by the branch-and-bound technique. This technique, and variations of it, has been applied with considerable success. Another approach is based on dynamic programming. The combination of dynamic programming with branch-and-bound techniques appears to be a promising area of research.

In the mid-1930's, Turing divided all imaginable problems in mathematics into two classes: (1) problems for which algorithms can be written at least in principle, and (2) problems for which algorithms can never be written. The first class is again subdivided into two groups: (1) problems which have polynomial-time algorithms, and (2) problems with exponential-time algorithms, i.e., the solution time for a network grows with its size relatively slowly like some polynomial, or explodes like some exponential function, respectively. For example, the well known Euler problems, which ask whether there is a path through a network (graph) that traverses each line exactly once, belong to the group with polynomial-time algorithms. On the other hand, Hamilton problems

(travelling salesman problems), which ask whether there is a path through a network that touches each node exactly once, belong to the group for which only exponential-time algorithms are known at this time.

Since many of the real world problems tend to have only exponential-time algorithms, they may not be susceptible to any of the analytical techniques of operations research. Alternative approaches are based on heuristic techniques. The phrase "rule of thumb" is often used synonymously with "heuristic." Heuristic techniques are strategies for seeking a method or methods which might produce a solution to a particular problem, although not necessarily the optimum solution. They involve the development of a set of heuristic rules, which hopefully will aid in the discovery of one or more satisfactory solutions to a problem.

In recent years, a good deal of work has been done in operations research in the development of heuristic programs for solving large combinatorial problems. Most of the interesting developments in heuristic programming described in the literature take the form of computer programs. A number of different heuristic programs for scheduling problems with limited resources have been developed in the past few years. Heuristic programs for resource scheduling usually take one of two forms: (1) resource leveling programs which attempt to reduce peak resource requirements within a constraint on the overall duration of activities, and (2) resource allocation programs which allocate available resources to individual activities in an attempt to find the shortest overall schedule consistent with fixed resource limits.

C. Artificial Intelligence

Artificial intelligence has had its problems with definitions. It appears that a completely satisfactory definition has not yet been found. Here is another unsatisfactory one: Artificial intelligence is a synthesis of computer activity elements which, as a whole, resembles human problem solving processes. The central goals of artificial intelligence research are to make computers more useful and to understand the principles which make intelligence possible.

By application areas, artificial intelligence comprises such topics as: (1) natural language processing, (2) expert consulting systems, (3) theorem proving, (4) combinatorial and scheduling problems, (5) perception problems, (6) automatic programming, (7) robotics, and (8) intelligent retrieval from databases. In the following parts of this section, these topics are briefly discussed with occasional reference to the presentation material in Volume II of this report. The mission operations model, shown in Figure 4, will be used to establish, by example, relevance to the space program and certain of its operational elements.

To be sure, the topics of artificial intelligence also relate to many other areas in the space program not depicted in Figure 4, e.g., computer-aided design, computer-aided manufacturing, computer-aided testing, computer-aided management, etc. In particular, they relate to the higher level decision-making and problem solving processes in an operational systems hierarchy as discussed above.

Natural Language Processing - If in Figure 4 the human operator of the system wants the machine to do something, he must be able to communicate his wishes to the machine. There are various possibilities for doing this. These are generally based on digital (tactile) input, or on analog input which is translated into digital data acceptable to the computer. In recent years, systems have been developed that recognize single spoken words with good reliability. The problem of recognizing connected speech has not been solved, except in very simple cases. This offers the possibility to communicate simple commands to the machine by voice. It is one more channel that supplements the digital channel and, in principle, is similar to the digital input. The machine recognizes an acoustic pattern but does not understand the language input in context.

Natural language understanding has been researched by AI investigators for years with modest success. It has been very difficult to develop computer systems capable of generating and "understanding" even fragments of a natural language, such as English. A computer system capable of understanding a message in natural language would require both the context and knowledge and the processes for making the inferences assumed by the message generator. Some progress has been made toward computer systems of this sort. Fundamental to the development of such systems are certain artificial intelligence ideas about the structures for representing contextual knowledge and certain techniques for making inferences from that knowledge.

There are now several programs that appear capable of distinguished dialogue on restricted subjects such as the toy-blocks world and Moon rocks. Recent advanced work is based on the idea of word expert parser, where each word

in the text has its own expert consultant who has expert knowledge about all the different meanings of his word. If a sentence is to be parsed, all the required expert consultants are called up. They communicate with each other independently and in clusters until the sentence is parsed, and the sentence is understood within a given context. The approach appears highly promising.

Expert Consulting Systems - The operation of large systems, such as a space mission, requires frequently more than one human operator. The Voyager space mission has had, at times, up to 250 people involved in the mission operations facility. A certain request for a command to the spacecraft usually originates with the science team. It then goes to the science integration team and to the sequence planning team. Finally, a command sequence is prepared by the sequence implementation group, ready to be sent to the spacecraft. During this process, extensive iterations are done between these teams of experts and other expert teams (structural engineers, control engineers, propulsion engineers, etc.) responsible for the spacecraft's welfare and navigation. These experts bring to bear their knowledge and experience to develop a command sequence which will satisfy the original request. The command sequence must have a consistent logical structure, and its execution must not violate system constraints. Automating this process, in whole or in part, requires so-called automatic expert consulting systems. Such an automatic expert consulting system incorporates all or part of the problem specific (space mission specific) knowledge embedded in the above mentioned expert teams.

Expert consulting systems provide human users with expert conclusions about specialized (domain specific) subject areas. Such systems have been built, and they can diagnose diseases, evaluate potential ore deposits, suggest

structures for complex organic chemicals, and even provide advice about how to use other computer systems. They can be used as experts in computer-aided design. For instance, programs have been developed capable of understanding electronic circuits. These programs reach conclusions about electronic circuits using human (engineering)-like reasoning and their explanations of complicated electronic devices are in terms that are easily understood by electrical engineers.

A key problem in the development of expert consulting systems is how to represent and use the knowledge that human experts in these subjects possess and use. This problem is made more difficult by the fact that the expert knowledge in many important fields is often imprecise, uncertain, or anecdotal. Many expert consulting systems employ the artificial intelligence technique of rule-based deduction. In such systems, expert knowledge is represented as a large set of simple rules, and these rules are used to guide the dialogue between the system and the user and to deduce conclusions.

Theorem Proving - The study of theorem proving has been of significant value in the development of artificial intelligence methods. Many informal tasks, including medical diagnosis, information retrieval, spacecraft command sequencing, etc., can be formalized as theorem-proving problems. The formalization of the deductive process in using the language of predicate logic, for example, helps to understand more clearly some of the components of reasoning. For instance, a skilled mathematician uses what he might call judgment to reduce the main problem into subproblems which can be worked on independently, and to guess successfully about which previously proven theorems in a subject area will be useful in a proof.

In its most elementary form, a representation for problem reduction must specify the decomposition of problems into subproblems. Planning representations more advanced than the classical AND-OR problem reduction graphs need to include other information in addition to subproblem decompositions. Among the additional information desired in a plan representation are (1) the relationships between actions and goals, (2) the side-effects or byproducts of actions, and (3) the interactions among goals. The Common Sense Algorithm is a network representation of functional relationships in systems or processes. It makes explicitly available to the problem solver the types of information enumerated above. This is accomplished by using network descriptions whose nodes (events) represent actions and states, and whose links represent a small set of allowable relations among them.

Combinatorial and Scheduling Problems - The scheduling and sequencing of activities on a spacecraft must take into consideration resource constraints in terms of time, power, propulsion, mass, etc. Such optimal scheduling problems are combinatorially expanding with some measure of problem size. A classical example is the travelling salesman's problem which has been mentioned above. Many of these problems can be and have been attacked by artificial intelligence methods. These efforts were directed at making the time-versus-problem-size curve grow as slowly as possible, even when it must grow exponentially. Several methods have been developed for delaying and moderating the inevitable combinatorial explosion. Again, knowledge about the problem domain is the key to more efficient solution methods. Many of the methods developed to deal with combinatorial problems are also useful on other, less combinatorially severe problems.

Perception Problems - The spacecraft or the remote robot (Fig. 4) is equipped with a sensory system which collects data about its surrounding environment and about its own state. Based on this sensor data, the robot's computer system should be able to make appropriate decisions autonomously. To do that, the sensor data must be perceived and "understood," and this understanding requires a large base of knowledge about the things being perceived.

The process of perception studied in artificial intelligence usually involves a set of operations. A visual scene, say, is encoded by sensors and represented as a matrix of intensity values. These are processed by detectors that search for primitive picture components such as line segments, simple curves, corners, etc. These, in turn, are processed to infer information about the three-dimensional character of the scene in terms of its surfaces and shapes. The ultimate goal is to represent the scene by some appropriate model which might consist of a high-level description.

The point of the whole perception process is to produce a condensed representation which can be substituted for the unmanageably immense raw input data. Obviously, the nature and quality of the final representation depend on the goals of the perceiving system. If colors are important, they must be noticed; if spatial relationships and measurements are important, they must be judged accurately. Different systems have different goals, but all must reduce the tremendous amount of sensory data at the input to a manageable and meaningful description. This reduction and representation formation process is a decision-making and problem solving process of the most intricate kind. Accurate decisions must be made about which data to retain, and then the problem needs to be solved of how the retained data is to be correctly synthesized to give a meaningful representation.

The strategy of making hypotheses about various levels of description and then testing these hypotheses seems to offer an approach to this problem. Systems have been constructed that process suitable representations of a scene to develop hypotheses about the components of a description. These hypotheses are then tested by detectors that are specialized to the component descriptions. The outcomes of these tests, in turn, are used to develop better hypotheses, etc. This hypothesize-and-test paradigm is applied at many levels of the perception process. Several aligned segments suggest a straight line; a line detector can be employed to test it. Adjacent rectangles suggest the faces of a solid prismatic object; an object detector can be employed to test it.

The process of hypothesis formation requires a large amount of knowledge about the expected scenes. Some Artificial Intelligence reseachers have suggested that this knowledge be organized in a special structure called a frame or schema. For example, when a robot approaches an antenna in space, it activates an antenna schema, which loads into a working memory a number of expectations about what might be seen next. Suppose the robot perceives a circular form. This form, in the context of an antenna schema, might suggest a circular feed. The feed schema might contain the knowledge that feeds typically do not touch the antenna surface. A special detector, applied to the scene, confirms this expectation, thus raising confidence in the antenna feed hypothesis.

Automatic Programming - The sensory data collected by the robot or the remote system (Fig. 4) and the understanding derived from the extracted information is used in the development of plans and decisions for actions by the

robot. Occasionally, this calls for actions that were not anticipated and were not preprogrammed. A new program must then be written, possibly automatically.

What is meant here by automatic programming might be described as a "super-compiler" or a program that could take in a very high-level description of what the program is to accomplish, and from it produce a program. The high-level description might be a precise statement in a formal language, such as the predicate calculus, or it might be a loose description, say, in English, that would require further dialogue between the system and the user in order to resolve ambiguities.

The task of writing a computer program is also related to other areas of artificial intelligence. Much of the basic research in automatic programming, theorem proving, and robot problem solving overlaps. In a sense, existing compilers already do "automatic programming." They take in a complete source code specification of what a program is to accomplish, and they write an object code program to do it.

The task of automatically writing a program to achieve a stated result is closely related to the task of proving that a given program achieves a stated result. The latter is called program verification. Many automatic programming systems produce a verification of the output program as an added benefit.

One of the important contributions of research in automatic programming has been the notion of debugging as a problem-solving strategy. It has been found that it is often much more efficient to produce an errorful solution to a programming or robot control problem cheaply and then modify it (to make it work correctly), than to insist on a first solution completely free of defects.

Robotics - The problem of controlling the physical actions of a mobile robot might not seem to require much intelligence. Even small children are able to navigate successfully through their environment and to manipulate items, such as light switches, toy blocks, eating utensils, etc. However, these same tasks, performed almost unconsciously by humans, when performed by a machine require many of the same abilities used in solving more intellectually demanding problems.

Research on robots or robotics during the last two decades has helped to develop and check out many artificial intelligence ideas in automated problem solving and decision making. It has led to several techniques for modeling world states and for describing the process of change from one world state to another. It has led to a better understanding of how to generate plans for action sequences and how to monitor the execution of these plans. Complex robot control problems have forced researchers to develop methods for planning first at a high level of abstraction, ignoring details, and then at lower and lower levels, where details become important.

Intelligent Retrieval from Databases - The data that are sent back by the remote system to the ground system (or control station, Fig. 4) are for many space missions immense. This is especially true for deep space missions and many global service missions. Therefore, large database systems are required.

The design of database systems is an active subspecialty of computer science, and many techniques have been developed to enable the efficient representation, storage, and retrieval of large numbers of facts. From an artificial

intelligence point of view, the subject becomes interesting when one wants to retrieve answers that require deductive reasoning with the facts in the database.

There are several problems that confront the designer of such an intelligent information retrieval system. First, there is the immense problem of building a system that can understand queries stated in a natural language like English. Second, even if the language-understanding problem is dodged by specifying some formal, machine-understandable query language, the problem remains of how to deduce answers from stored facts. Third, understanding the query and deducing an answer may require knowledge beyond that explicitly represented in the subject domain database. Common knowledge (typically omitted in the subject domain database) is often required.

The data management problem grows as the number of users and the number of data bases increases. One suggestion to relieve the problem is to create knowledge base centers to service a large number of users. Knowledge base centers would retrieve requests and transmit responses to users who access the centers through remote terminals.

Knowlege base centers also must be able to solve computational problems at request. Considerable research has been done on the automation of such distributed problem solving systems, and several distributed problem solving systems have been constructed or proposed: (1) HEARSAY II, (2) distributed HEARSAY II architecture, (3) HARPY machine, (4) distributed NOAH, (5) traffic control, and (6) contract nets. However, few truly distributed problem solving systems exist. The prime issues in designing tools for distributed problem solving are

(1) the role of intercommunication is unclear, (2) the problems of addressing and binding of messages are unresolved, (3) the handling of asynchronous events is still a problem, and (4) many questions of fault tolerance need to be answered.

D. The Human Problem Solver

The involvement of the human operator in an operational system (Fig. 4) depends on the level of the automatic decision-making and problem solving capability built into the supporting computers, at the remote system as well as at the ground system. Well-structured problems can usually be programmed to be solved by the computer, while ill-structured problems are still reserved for the human problem solver, although certain subproblems may be computer programmed in support of the human decision-maker. Appropriate design of the system architecture may facilitate nonprogrammed (ill-structured) as well as programmed (well-structured) decision-making. Programmed activity tends to drive out non-programmed activity as understanding of the problems increases in terms of computer programming. The frontier is progressively moving towards increased automation, and the human decision-maker becomes more of a supervisor. This has led to the term "supervisory control" for teleoperator systems. It has also led to an area of research that has to do with the appropriate (optimum) allocations of functions between man and machine.

Human information processing is the study of the psychological mechanisms underlying mental functioning. Memory, problem solving, language, perception, thinking -- these are some of the major areas studied. In the past decade there have been sufficient systematic advances in our knowledge that these areas now constitute perhaps the best understood problems in contemporary psychology.

Studies regarding human attention are of special importance to problems faced by NASA. Humans have limited mental resources, and the deployment of these resources constitutes an important part of behavior. These limitations appear to apply primarily to conscious control.

Tasks that require conscious decision-making or control can suffer in times of stress or when other tasks must be performed or thought about simultaneously. When several tasks simultaneously demand a share of conscious resources, deterioration of performance results. Tasks that are learned well enough that they appear "automated" seem to suffer little as a result of other activities.

Almost all knowledge regarding human performance deals with processing or arriving information, or the operation of the human as an element of a control structure. Not enough is known about the nature of conscious and subconscious control mechanisms, and not enough is known about the various modes of operation of the human. Little is known about the human's ability to interact with and control the environment, particularly under the unique conditions faced in NASA programs.

This area of knowledge about the human has many potential applications for NASA. On the spacecraft, at mission control, onsite during a mission, all these situations require different aspects of human capability. They are critical to the success of NASA's missions, the more so as missions become longer, more complex, with space repair, manufacture, and mining as possible tasks.

VI. CONCLUSIONS AND RECOMMENDATIONS

Automated decision-making and problem solving were of prime concern at this conference with the aim to explore commonalities and differences among the sub-areas: advanced control theory, operations research, and artificial intelligence. Existing techniques were assessed, trends of developments were determined, and some potential applications in the NASA space program were identified. This was accomplished at varying depths of penetration into the technical details of each area. Volume II of this report contains the material presented at the Conference. Unfortunately, not all interesting and relevant topics could be covered at an interdisciplinary conference of this type (or at any conference). However, it is felt that the areas that were covered and discussed are indicative of the advances that have been and need to be made, and that there was keen awareness and occasional appreciation for the problems of another discipline. Specifically, the topics that were described in each disciplinary area as being at the edge of technological development and in need of vigorous research are identified in the following.

<u>Advanced Control Theory</u> - Hierarchically, intelligent man-machine interactive systems should be developed as a natural successor to adaptive learning and self-organizing control systems. In the analytical area, it is required (1) to develop new techniques for decomposing and simplifying problems, (2) to develop approaches to special classes of decentralization problems to handle second guessing, signalling and protocols, and (3) to develop new classes of models which are susceptible to analysis.

Operations Research - Modeling and analysis techniques need to be developed that deal with multi-objective criteria, system complexity, and heuristic approaches. The combination of dynamic programming and branch-and-bound techniques should be explored. New advances in graph theory should be exploited for large scale systems modeling, and computer-aided interactive modeling and analysis for decision-making should be pursued.

Artificial Intelligence - Models of: (1) natural language understanding, (2) human understanding of causality, and (3) general knowledge storage are required. Expert systems need to be developed which are capable of expert-level performance in relevant areas with domain specific knowledge that is represented naturally and that is used in an understandable line of reasoning. Artificial intelligence techniques should be incorporated into computer-aided design, computer-aided testing, and computer-aided planning approaches. Good real-time display oriented human interfaces are required.

From this it can be concluded that there are considerable overlaps of technical concerns. The Venn diagram in Figure 7 is indicative of the kinds of common interests among advanced control theory, operations research, and artificial intelligence. In these common areas, it is not only possible that these three disciplines can learn from each other, but the solution of some problems may not be possible without the tools and the way of looking at things of the other disciplines.

To facilitate intercommunication, interaction, and cooperation among diverse groups or disciplines, each with its own "culture," is at best difficult.

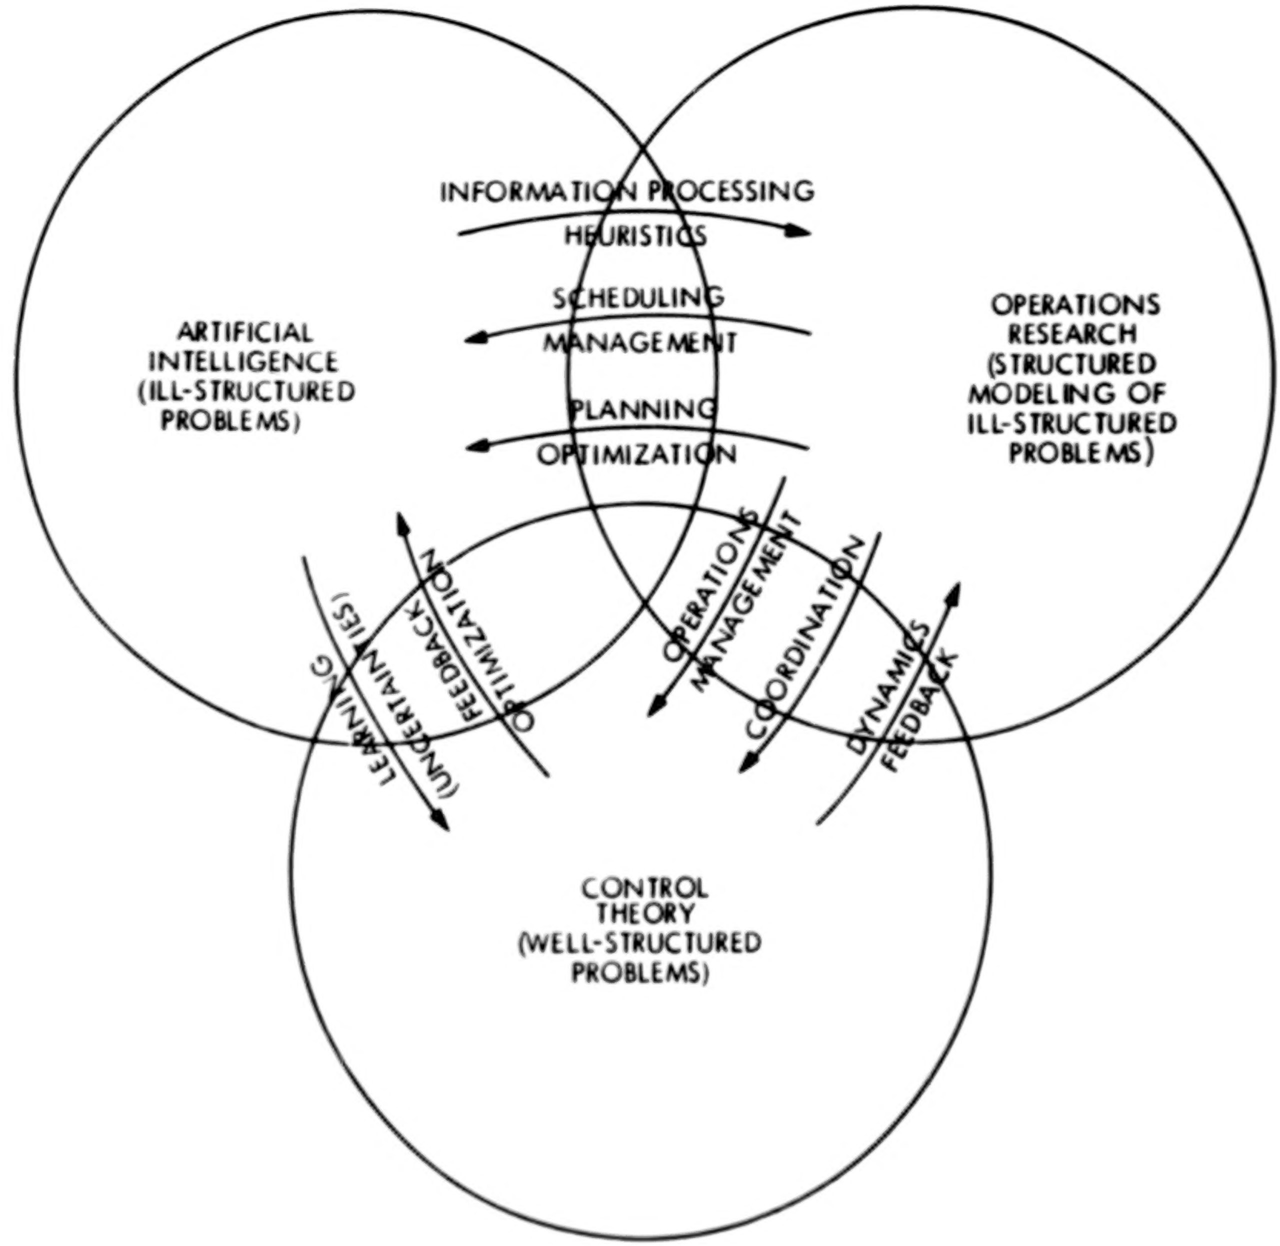

Figure 7. Venn Diagram of Interdisciplinary Issues and Expertise

What is required is a common purpose for doing so. A focus in the form of a project is needed. The implementation of such a project should require the involvement of the three disciplines at approximately the same level of effort in terms of resources and complexity. It should be exciting enough to capture the imagination of the researchers (and their students). It should be narrow enough to fill specific NASA needs, it should be broad enough so that the research for it would also fill needs even if NASA were not involved, and it should be about ten to fifteen years in the future, so that fundamental research would still have a chance to contribute to and impact the project. Such a project with specific research problems would promote an organization along defined lines, and people would tend to organize themselves along these lines.

An example of such a project would be a space manufacturing facility to be in operation in about 15 years. This space station would contain several different kinds of teleoperators along with many other kinds of manufacturing equipment. About 5 to 10 people would man the facility and would control it from a command station of some type. The structure would be very light and flexible, thus requiring complex systems to control and stabilize it. The movements of masses, teleoperators and equipment would require an intricate scheduling and resource management system, and the required robotics and teleoperator capabilities would put high demands on artificial intelligence techniques and on man-machine function allocation trade-offs.

REFERENCES

1. Final Report of the NASA Study Group on Machine Intelligence and Robotics, NASA, Jet Propulsion Laboratory, Report No. 715-32, March 1980.

2. "Space Operations Center - A Concept Analysis," NASA, Johnson Space Center, Report No. JSC-16277, November 29, 1979.

3. "The New Science of Management Decision," Revised Edition, by H. A. Simon, Prentice-Hall, Inc. 1977.

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA CP-2180		
4. Title and Subtitle AUTOMATED DECISION MAKING AND PROBLEM SOLVING Volume I. Executive Summary		5. Report Date August 1981
		6. Performing Organization Code 506-54-73-01-00
7. Author(s) Ewald Heer		8. Performing Organization Report No. L-14396
9. Performing Organization Name and Address NASA Langley Research Center Hampton, VA 23665		10. Work Unit No.
		11. Contract or Grant No.
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546		13. Type of Report and Period Covered Conference Publication
		14. Sponsoring Agency Code

15. Supplementary Notes
Ewald Heer: University of Southern California, Los Angeles, California.

16. Abstract

On May 19-21, 1980, NASA Langley Research Center hosted a Conference on Automated Decision Making and Problem Solving. The purpose of the conference was to explore related topics in artificial intelligence, operations research, and control theory and, in particular, to assess existing techniques, determine trends of development, and identify potential for application in automation technology programs at NASA. The first two days consisted of formal presentations by experts in the three disciplines. The third day was a workshop in which the invited speakers and NASA personnel discussed current technology in automation and how NASA can and should interface with the academic community to advance this technology. The conference proceedings are published in two volumes: Volume I is an executive summary; Volume II contains the papers presented at the conference.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
Automation, Decision making, Problem solving, Artificial intelligence, Operations research, Control theory	Unclassified - Unlimited Subject Category 66

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	55	A04

For sale by the National Technical Information Service, Springfield, Virginia 22161

NASA-Langley, 1981

90%

END

3-23-82